# Joint Attention in Autonomous Driving (JAAD)

Iuliia Kotseruba, Amir Rasouli and John K. Tsotsos
{yulia_k, aras, tsotsos}@cse.yorku.ca

arXiv:1609.04741v5 [cs.RO] 24 Apr 2017

***Abstract*—In this paper we present a novel dataset for a critical aspect of autonomous driving, the joint attention that must occur between drivers and of pedestrians, cyclists or other drivers. This dataset is produced with the intention of demonstrating the behavioral variability of traffic participants. We also show how visual complexity of the behaviors and scene understanding is affected by various factors such as different weather conditions, geographical locations, traffic and demographics of the people involved. The ground truth data conveys information regarding the location of participants (bounding boxes), the physical conditions (e.g. lighting and speed) and the behavior of the parties involved.**

## I. Introduction

Autonomous driving has been a topic of interest for decades. Implementing autonomous vehicles can have a great economic and social impacts including reducing the cost of driving, increasing fuel efficiency and safety, enabling transportation for non-drivers and reducing the stress of driving by allowing motorists to rest and work while traveling [1]. As for the macroeconomic impacts, it is estimated that autonomous vehicles industry and related software and hardware technologies will account for a market size of more than 40 billion dollars by 2030 [2].

Partial autonomy has long been used in commercial vehicles in the form of technologies such as cruise control, park assist, automatic braking, etc. Fully autonomous vehicles also have been successfully developed and tested under certain conditions. For example, the DARPA challenge 2005 set the task of autonomously driving a 7.32 miles predefined terrain in the deserts of Nevada. Out of 23 final contestants 4 cars successfully completed the course within the allowable time limit (10 hours) while driving fully autonomously [3].

Despite such success stories in autonomous control systems, designing fully autonomous vehicles for urban environments still remains an unsolved problem. Aside from challenges associated with developing suitable infrastructures and regulating the autonomous behaviors [1], in order to be usable in urban environments autonomous cars must have a high level of precision and meet very high safety standards [4].

Today one of the major dilemmas faced by autonomous vehicles is how to interact with the environment including infrastructure, cars, drivers or pedestrians [5], [6], [7]. The lapses in communication can be a source of numerous erroneous behaviors [8] such as failure to predict the movement of other vehicles [9], [10] or to respond to unexpected behaviors of other drivers [11].

The impact of perceptual failures on the behavior of an autonomous car is also evident in the 2015 annual report on Google's self-driving car [12]. This report is based on testing self-driving cars for more than $424,000$ miles of driving on public roads including both highways and streets. Throughout these trials, a total of 341 disengagements occurred in which the driver had to take over the car, and about $90\%$ of the cases occurred in busy streets. The interesting implication here is that over $1/3$ of the disengagements were due to"perception discrepancy" in which the vehicle was unable to understand its environment and about 10% of the cases were due to incorrect prediction of traffic participants and inability to respond to reckless behaviors.

There have been a number of recent developments to address these issues. A natural solution is establishing wireless communication between traffic participants. This approach has been tested for a number of years using cellular technology [7], [13]. This technique enables vehicle to vehicle (V2V) and vehicle to infrastructure (V2I) communication allowing tasks such as Cooperative Adaptive Cruise Control (CACC), improving positioning technologies such as GPS, and intelligent speed adoption in various roads. Peer to peer traffic communication is expected to enter the market by 2019.

Although V2V and V2I communications are deemed to solve a number of issues in autonomous driving, they also have a number of drawbacks. This technology relies heavily on cellular technology which is costly and has much lower reliability compared to traditional sensors such as radars and cameras. In addition, communication highly depends on all parties functioning properly. A malfunction in any communication device in any of the systems involved can lead to catastrophic safety issues.

Maintaining communication with pedestrians is even more important for safe autonomous driving. Inattention from both drivers and pedestrians regardless of their knowledge of the vehicle code is one of the major reasons for traffic accidents, most of which involve pedestrians at crosswalk locations [14].

Honda recently released a new technology similar to V2V communication that attempts to establish a connection with pedestrians through their cellular phones [15]. Using this method, the pedestrian's phone broadcasts its position warning the autonomous car that a pedestrian is about to cross the street so the car can respond accordingly. This technology also can go one step further and inform the car about the state of the pedestrian, for instance, whether he/she is listening to music, texting or is on a call. Given the technological and regulatory obstacles to developing such technologies, using them in the near future does not seem feasible.

In late 2015 Google patented a different technology to communicate with pedestrians using a visual interface called pedestrian notifications [16]. In this approach, the Google car estimates the trajectories of pedestrian movements. If the car finds the behavior of a pedestrian to be uncertain (i.e. cannot decide whether or not he/she is crossing the street), it notifies

the corresponding pedestrian about the action it is about to take using a screen installed on the front hood of the car. Another proposed option for communication is via a sound device or other kinds of physical devices (possibly a robotic arm). This technology has been criticized for being distracting and lacking the ability to efficiently communicate if more than one pedestrian is involved.

Given the problems associated with establishing explicit communication with other vehicles and pedestrians, Nissan, in their latest development, announced a passive method of dealing with uncertainties in the behavior of pedestrians [17], [18], in attempt to understand human drivers and pedestrians behaviors in various traffic scenarios. The main objective of this work is to passively predict pedestrian behavior using visual input and only use an interface, e.g. a green light, to inform them about the intention of the autonomous car.

Toyota in partnership with MIT and Stanford recently announced using a similar passive approach toward autonomous driving [19]. Information such as the type of equipment the pedestrian carries, his/her pose, direction of motion and behavior as well as the human driver's reactions to events are extracted from videos of traffic situations and their 3D reconstructions. This information is used to design a control system for determining what course of action to take in given situations. At present no data resulting from this study has been released and very little information is available about the type of autonomous behavior the scientists are seeking to obtain.

Given the importance of pedestrian safety, multiple studies were conducted in the last several decades to find factors that influence decisions and behavior of traffic participants. For example, a driver's pedestrian awareness can be measured based on whether the driver is decelerating upon seeing a pedestrian ([20], [21], [22]). Several recent studies also point out that pedestrian's behavior, such as establishing eye-contact, smiling and other forms of non-verbal communication, can have a significant impact on the driver's actions ([23], [24]). Although a majority of these studies are aimed at developing better infrastructure and traffic regulations, their conclusions are relevant for autonomous driving as well.

In an attempt to better understand the problem of vehicle to vehicle (V2V) and vehicle to pedestrian (V2P) communication in the autonomous driving context we suggest viewing it as an instance of joint attention and discuss why existing approaches may not be adequate in this context. We propose a novel dataset that highlights the visual and behavioral complexity of traffic scene understanding and is potentially valuable for studying the joint attention issues.

## II. Autonomous driving and joint attention

According to a common definition, joint attention is the ability to detect and influence an observable attentional behavior of another agent in social interaction and acknowledge them as an intentional agent [25]. However, it is important to note that joint attention is more than simultaneous looking, attention detection and social coordination, but also includes an intentional understanding of the observed behavior of others.

Since joint attention is a prerequisite for efficient communication, it has been gaining increasing interest in the fields of robotics and human-robot interaction. Kismet [26] and Cog [27], both built at MIT in the late 1990s, were some of the first successes in social robotics. These robots were able to maintain and follow eye gaze, reacted to the behavior of their caregivers and recognized simple gestures such as declarative pointing. More recent work in this area is likewise concerned with gaze following [28], [29], [27], pointing [30], [27] and reaching [31], turn-taking [32] and social referencing [33]. With a few exceptions [34], [35], almost all joint attention scenarios are implemented with stationary robots or robotic heads according to a recent comprehensive survey [36].

Surprisingly, despite increasing interest for joint attention in the fields of robotics and human-robot interaction, it has not been explicitly mentioned in the context of autonomous driving. For example, communication between the driver and pedestrian is an instance of joint attention. Consider the following scenario: a pedestrian crossing the street (shown in Figure 1a). Initially she is looking at her cell phone, but as she approaches the curb, she looks up and slows down because the vehicle is still moving. When the car slows down, she speeds up and crosses the street. In this scenario all elements of joint attention are apparent. Looking at the car and walking slower is an observable attention behavior. The driver slowing down the car indicates that he noticed the pedestrian and is yielding. His intention is clearly interpreted as such, as the pedestrian speeds up and continues to cross. A similar scene is shown in Figure 1d. Here the pedestrian is standing at the crossing and looking both ways to find a gap in traffic. Again, once she notices that the driver is slowing down, she begins to cross.

While these are fairly typical behaviors for marked crossings, there are many more possible scenarios of communication between the traffic participants. Humans recognize a myriad of "social cues" in everyday traffic situations. Apart from establishing eye contact or waving hands, people may be making assumptions about the way a driver would behave based on visual characteristics such as the car's make and model [6]. Understanding these social cues is not always straightforward. Aside from visual processing challenges such as variation in lighting conditions, weather or scene clutter, there is also a need to understand the context in which the social cue is observed. For instance, if the autonomous car sees someone waving his hand, it needs to know whether it is a policeman directing traffic, a pedestrian attempting to cross the street or someone hailing a taxi. Consider Figure 1b, where a man is crossing the street and makes a slight gesture as a signal of yielding to the driver, or Figure 1c, where a man is jaywalking and acknowledges the driver with a hand gesture. Responding to each of these scenarios from the driver's perspective can be quite different and would require high-level reasoning and deep scene analysis.

Today, automotive industry giants such as BMW, Tesla, Ford and Volkswagen, who are actively working on autonomous driving systems, rely on visual analysis technologies developed by Mobileye[1] to handle obstacle avoidance, pedestrian

[1] http://www.mobileye.com/

(a)

(b)

(c)

(d)

Figure 1: Examples of joint attention

detection or traffic scene understanding. Mobileye's approach to solving visual tasks is to use deep learning techniques which require a large amount of data collected from hundreds of hours of driving. This system has been successfully tested and is currently being used in semi-autonomous vehicles. However, the question remains open whether deep learning suffices for achieving full autonomy in which tasks are not limited to detection of pedestrians, cars or obstacles (which are not still fully reliable [37], [38]), but also involve merging with ongoing traffic, dealing with unexpected behaviors such as jaywalking, responding to emergency vehicles, and yielding to other vehicles or pedestrians at intersections.

To answer this question we need to consider the following characteristics of deep learning algorithms. First, even though deep learning algorithms perform very well in tasks such as object recognition, they lack the ability to establish causal relationships between what is observed and the context in which it has occurred [39], [40]. This problem also has been empirically demonstrated by training neural networks over various types of data [39].

The second limitation of deep learning is the lack of robustness to changes in visual input [41]. This problem can occur when a deep neural network misclassifies an object due to minor changes (at a pixel level) to an image [42] or even recognizes an object from a randomly generated image [43].

## III. Existing datasets

The autonomous driving datasets currently available to the public are primarily intended for applications such as 3D mapping, navigation, and car and pedestrian detection. Out of these datasets only a limited number contain data that can be used for behavioral studies. Below some of these datasets are listed.

- KITTI [44]: This is perhaps one of the most known publicly available datasets for autonomous driving. It contains data collected from various locations such as residential areas and city streets, highways and gated environments. The main application is for 3D reconstruction, 3D detection, tracking and visual odometry. Some of the videos in KITTI show pedestrians, other vehicles and cyclists movements alongside the car. The data has no annotation of their behaviors.
- Caltech pedestrian detection benchmark [45]: This is a very large dataset of pedestrians consisting of approximately 10 hours of driving in regular traffic in urban environments. The annotations include temporal correspondence between bounding boxes around pedestrians and detailed occlusion lables.
- Berkeley pedestrian dataset [46]: This dataset consists of a large number of videos of pedestrians collected from a stationary car at street intersections. Bounding boxes around pedestrians are provided for pedestrian detection and tracking.
- Semantic Structure From Motion (SSFM) [47]: As the name implies, this dataset is collected for scene understanding. The annotation is limited to bounding boxes around the objects of interest and name tags for the purpose of detection. This dataset includes a number of street view videos of cars and pedestrians walking.
- The German Traffic Sign Detection Benchmark [48]: This dataset consists of 900 high-resolution images of roads and streets some of which show pedestrians crossing and cars. The ground truth for the dataset only specifies the positions of traffic signs in the images.
- The .enpeda.. (environment perception and driver assistance) Image Sequence Analysis Test Site (EISATS) [49]: EISATS contains short synthetic and real videos of cars driving on roads and streets. The sole purpose of this dataset is comparative performance evaluation of stereo vision and motion analysis. The available annotation is limited to the camera's intrinsic parameters.
- Daimler Pedestrian Benchmark Datasets [50]: These are particularly useful datasets for various scenarios of pedestrian detection such as segmentation, classification, and path prediction. The sensors of choice are monocular and binocular cameras and the datasets contain both color and grayscale images. The ground truth data is limited to the detection applications and does not include any behavioral analysis.
- UvA Person Tracking from Overlapping Cameras Datasets [51]: These datasets mainly are concerned with the tasks of tracking, and pose and trajectory estimation using multiple cameras. The ground truth is also limited to only facilitate tracking applications.

In recent years traffic behavior of drivers and pedestrians has became a widely studied topic for collision prevention and traffic safety. Several large-scale naturalistic driving studies have been conducted in the USA [52], [53], [54], which accumulated over 4 petabytes of data (video, audio, instrumental, traffic, weather, etc) from hundreds of volunteer drivers in multiple locations. However, only some depersonalized general statistics are available to the general public [55], while only qualified researchers have access to the raw video and sensor data.

## IV. The JAAD Dataset

The JAAD dataset was created to facilitate studying the behavior of traffic participants. The data consists of 346 high-resolution video clips (5-15s) with annotations showing various situations typical for urban driving. These clips were extracted from approx. 240 hours of driving videos collected in several locations. Two vehicles equipped with wide-angle video cameras were used for data collection (Table I). Cameras were mounted inside the cars in the center of the windshield below the rear view mirror.

The video clips represent a wide variety of scenarios involving pedestrians and other drivers. Most of the data is collected in urban areas (downtown and suburban), only a few clips are filmed in rural locations. Many of the situations resemble the ones we have described earlier, where pedestrians wait at the designated crossings. In other samples pedestrians may be walking along the road and look back to see if there is a gap in traffic (Figure 4c), peek from behind the obstacle to see if it is safe to cross Figure 4d, waiting to cross on a

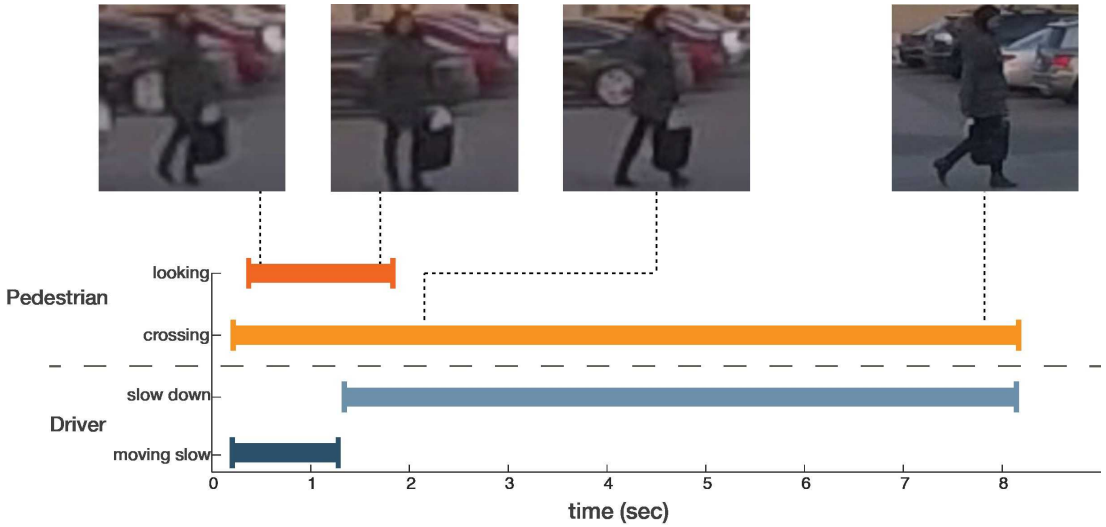


Figure 2: A timeline of events recovered from the behavioral data. Here a single pedestrian is crossing the parking lot. Initially the driver is moving slow and, as he notices the pedestrian ahead, slows down to let her pass. At the same time the pedestrian crosses without looking first, then turns to check if the road is safe, and as she sees the driver yielding, continues to cross. The difference in resolution between the images is due to the changes in distance to the pedestrian as the car moves forward.

divider between the lanes, carrying heavy objects or walking with children or pets. Our dataset captures pedestrians of various ages walking alone and in groups, which may be a factor affecting their behavior. For example, elderly people and parents with children may walk slower and be more cautious.

| # of clips | Location | Resolution | Camera model |
|---|---|---|---|
| 55 | North York, ON, Canada | 1920 × 1080 | GoPro HERO+ |
| 287 | Kremenchuk, Ukraine | 1280 × 720 | Highscreen Black Box Connect |
| 6 | Hamburg, Germany | 1280 × 720 | Highscreen Black Box Connect |
| 5 | New York, USA | 1920 × 1080 | GoPro HERO+ |
| 4 | Lviv, Ukraine | 1920 × 1080 | Garmin GDR-35 |

Table I: Locations and equipment used to capture videos in the JAAD dataset

The dataset contains fewer clips of interactions with other drivers, most of them occur in uncontrolled intersections, in parking lots or when another driver is moving across several lanes to make a turn.

Most of the videos in the dataset were recorded during the daytime and only a few clips were filmed at night, sunset and sunrise. The last two conditions are particularly challenging, as the sun is glaring directly into the camera (Figure 5b). We also tried to capture a variety of weather conditions (Figure 5), as yet another factor affecting the behavior of traffic participants. For example, during the heavy snow or rain people wearing hooded jackets or carrying umbrellas may have limited visibility of the road. Since their faces are obstructed it is also harder to tell if they are paying attention to the traffic from the driver's perspective.

We attempted to capture all of these conditions for further analysis by providing two kinds of annotations for the data: bounding boxes and textual annotations. Bounding boxes are provided only for cars and pedestrians that interact with or require attention of the driver (e.g. another car yielding to the driver, pedestrian waiting to cross the street, etc.). Bounding boxes for each video are written into an xml file with frame number, coordinates, width, height and occlusion flag.

Textual annotations are created using BORIS[2] [56] - event logging software for video observations. It allows to assign predefined behaviors to different subjects seen in the video, and can also save some additional data, such as video file id, location where the observation has been made, etc.

A list of all behaviors, independent variables and their values is shown in Table II. We save the following data for each video clip: weather, time of the day, age and gender of the pedestrians, location and whether it is a designated crosswalk. Each pedestrian is assigned a label (pedestrian1, pedestrian2, etc.). We also distinguish between the driver inside the car and other drivers, which are labeled as "Driver" and "Driver_other" respectively. This is necessary for the situations where two or more drivers are interacting. Finally, a range of behaviors is defined for drivers and pedestrians: walking, standing, looking, moving, etc.

An example of textual annotation is shown in Figure 3. The sequence of events recovered from this data is shown in Figure 2.

The dataset is available to download at http://data.nvision2.eecs.yorku.ca/JAAD_dataset.

## V. Conclusion

In this paper we presented a new dataset for the purpose of studying joint attention in the context of autonomous driving. Two types of annotations accompanying each video clip in the dataset make it suitable for pedestrian and car detection, as well as other areas of research, which could benefit from studying joint attention and human non-verbal communication, such as social robotics.

[2]http://www.boris.unito.it/

| Categorical variable | Values |
|---|---|
| time_of_day | day/night |
| weather | clear/snow/rain/cloudy |
| location | street/indoor/parking_lot |
| designated_crossing | yes/no |
| age_gender | Child/Young/Adult/Senior Male/Female |

| Behavior event | Type |
|---|---|
| Crossing | state |
| Stopped | state |
| Moving fast | state |
| Moving slow | state |
| Speed up | state |
| Slow down | state |
| Clear path | state |
| Looking | state |
| Look | point |
| Signal | point |
| Handwave | point |

Table II: Variables associated with each video and types of events represented in the dataset. There are two types of behavior events: state and point. State event may have an arbitrary duration, while point events last a short fixed amount of time (0.1 sec) and signify a quick glance or gestures made by pedestrians.

Observation id GOPR0103_528_542
Media file(s)
Player #1 GOPR0103_528_542.MP4
Observation date 2016-07-15 15:15:38
Description
Time offset (s) 0.000

Independent variables

| variable | value |
|---|---|
| weather | rain |
| age_gender | AF |
| designated | no |
| location | plaza |
| time_of_day | daytime |

| Time | Media file path | Media total length | FPS | Subject | Behavior | Comment | Status |
|---|---|---|---|---|---|---|---|
| 0.19 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | Driver | moving slow | | START |
| 0.208 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | pedestrian | crossing | | START |
| 0.308 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | pedestrian | looking | | START |
| 1.301 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | Driver | moving slow | | STOP |
| 1.302 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | Driver | slow down | | START |
| 1.892 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | pedestrian | looking | | STOP |
| 8.351 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | pedestrian | crossing | | STOP |
| 8.99 | GOPR0088_335_344.MP4 | 9.01 | 29.97 | Driver | slow down | | STOP |

Figure 3: Example of textual annotation for a video created using BORIS. The file contains the id and the name of the video file, a tab-separated list of independent variables (weather, age and gender of pedestrians, whether the crossing is designated or not, location and time of the day) and a tab-separated list of events. Each event has an associated time stamp, subject, behavior and status, which may be used to recover sequence of events for analysis.

## ACKNOWLEDGMENT

We thank Mr. Viktor Kotseruba for assistance with processing videos for this dataset.

## REFERENCES

[1] T. Litman. (2014, dec) Autonomous Vehicle Implementation Predictions: Implications for Transport Planning. Online. [Online]. Available: http://www.vtpi.org/avip.pdf

[2] (2014, nov) Think Act: Autonomous Driving. Online. [Online]. Available: https://new.rolandberger.com/wp-content/uploads/Roland{_}Berger{_}Autonomous-Driving1.pdf

[3] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, J. G. Philip Fong, G. H. Morgan Halpenny, K. Lau, M. P. Celia Oakley, V. Pratt, and P. Stang, "Stanley: The Robot that Won the DARPA Grand Challenge," *Journal of Field Robotics*, vol. 23, no. 9, pp. 661–692, 2006.

[4] N. Kalra and S. M. Paddock. (2006, apr) Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability. [Online]. Available: http://www.rand.org/pubs/research{_}reports/RR1478.html

[5] W. Knight. (2015, dec) Can This Man Make AI More Human? Online. [Online]. Available: https://www.technologyreview.com/s/544606/can-this-man-make-aimore-human/

[6] L. Gomes. (2014, jul) Urban Jungle a Tough Challenge for Google's Autonomous Cars. Online. [Online]. Available: https://www.technologyreview.com/s/529466/urban-jungle-a-tough-challenge-for-googles-autonomous-cars/

[7] G. Silberg and R. Wallace. (2012) Self-driving cars: The next revolution. Online. [Online]. Available: https://www.kpmg.com/Ca/en/IssuesAndInsights/ArticlesPublications/Documents/self-driving-cars-next-revolution.pdf

[8] S. E. Anthony. (2016, mar) The Trollable Self-Driving Car. Online. [Online]. Available: http://www.slate.com/articles/technology/future{_}tense/2016/03/

[9] M. Richtel and C. Dougherty. (2016, sep) Google's Driverless Cars Run Into Problem: Cars With Drivers. Online. [Online]. Available: http://www.nytimes.com/2015/09/02/technology/personaltech/google-says-its-not-the-driverless-cars-fault-its-other-drivers.html?{_}r=1

[10] (2016, feb) Google Self-Driving Car Project Monthly Report. Online. [Online]. Available: https://static.googleusercontent.com/selfdrivingcar/files/reports/report-0216.pdf

[11] W. Knight. (2013, oct) Driverless Cars Are Further Away Than You Think. Online. [Online]. Available: https://www.technologyreview.com/

(a)

(b)

(c)

(d)

Figure 4: More examples of joint attention

(a) Sunset

(b) Sunrise

(c) After a heavy snowfall

(d) During a heavy rain

(e) Multiple pedestrians crossing

(f) At the parking lot

Figure 5: Sample frames from the dataset showing different weather conditions and locations.

s/520431/driverless-cars-are-further-away-than-you-think/

[12] (2015, dec) Google Self-Driving Car Testing Report on Disengagements of Autonomous Mode. Online. [Online]. Available: https://static.googleusercontent.com/media/www.google.com/en//selfdrivingcar/files/reports/report-annual-15.pdf

[13] W. Knight. (2015) Car-to-Car Communication. Online. [Online]. Available: https://www.technologyreview.com/s/534981/car-to-car-communication/

[14] D. R. Ragland and M. F. Mitman, "Driver/Pedestrian Understanding and Behavior at Marked and Unmarked Crosswalks," Safe Transportation Research & Education Center, Institute of Transportation Studies (UCB), UC Berkeley, Tech. Rep., 2007. [Online]. Available: http://escholarship.org/uc/item/1h52s226

[15] (2016, may) Honda tech warns drivers of pedestrian presence. Online. [Online]. Available: http://www.cnet.com/roadshow/news/nikola-motor-company-the-ev-startup-with-the-worst-most-obvious-name-ever/

[16] C. P. Urmson, I. J. Mahon, D. A. Dolgov, and J. Zhu, "Pedestrian notifications," 2015. [Online]. Available: https://www.google.com/patents/US8954252

[17] G. Stern. (2015, feb) Robot Cars and Coordinated Chaos. Online. [Online]. Available: http://www.wsj.com/articles/robot-cars-and-the-language-of-coordinated-chaos-1423540869

[18] R. Jones. (2016, apr) T3 Interview: Nissan's research chief talks autonomous vehicles and gunning it in his Nissan GT-R Black Edition. Online. [Online]. Available: http://www.t3.com/features/t3-interview-dr-maarten-sierhuis-nissan-s-director-of-research-at-silicon-valley-talks-autonomous-vehicles-and-gunning-it-in-his-gt-r-black-edition

[19] E. Ackerman and E. Guizzo. (2015, sep) Toyota Announces Major Push Into AI and Robotics, Wants Cars That Never Crash. Online. [Online]. Available: http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/toyota-announces-major-push-into-ai-and-robotics

[20] M. Akamatsu, Y. Sakaguchi, and M. Okuwa, "Modeling of Driving Behavior When Approaching an Intersection Based on Measured Behavioral Data on An Actual Road," in *Human factors and Ergonomics*

Figure 6: A selection of images of pedestrians from the dataset

*Society Annual Meeting Proceedings*, vol. 47, no. 16, 2003, pp. 1895–1899.

[21] Y. Fukagawa and K. Yamada, "Estimating driver awareness of pedestrians from driving behavior based on a probabilistic model," in *IEEE Intelligent Vehicles Symposium (IV)*, 2013.

[22] M. T. Phan, I. Thouvenin, V. Fremont, and V. Cherfaoui, "Estimating Driver Unawareness of Pedestrian Based On Visual Behaviors and Driving Behaviors," in *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications*, 2014.

[23] Z. Ren, X. Jiang, and W. Wang, "Analysis of the Influence of Pedestrian's eye Contact on Driver's Comfort Boundary During the Crossing Conflict," *Green Intelligent Transportation System and Safety, Procedia Engineering*, 2016.

[24] N. Gueguen, C. Eyssartier, and S. Meineri, "A pedestrian's smile and driver's behavior: When a smile increases careful driving," *Journal of Safety Research*, vol. 56, pp. 83–88, 2016.

[25] F. Kaplan and V. V. Hafner, "The Challenges of Joint Attention," *Interaction Studies*, vol. 7, no. 2, pp. 135–169, 2006.

[26] C. Breazeal and B. Scassellati, "How to build robots that make friends and influence people," *Intelligent Robots and Systems*, pp. 858–863, 1999. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs{_}all.jsp?arnumber=812787

[27] B. Scassellati, "Imitation and Mechanisms of Joint Attention: A Developmental Structure for Building Social Skills on a Humanoid Robot," *Lecture Notes in Computer Science*, vol. 1562, pp. 176–195, 1999. [Online]. Available: http://www.springerlink.com/content/wljp04e4h5b4lthh/

[28] A. P. Shon, D. B. Grimes, C. L. Baker, M. W. Hoffman, Z. Shengli, and R. P. N. Rao, "Probabilistic gaze imitation and saliency learning in a robotic head," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2005, pp. 2865–2870, 2005.

[29] I. Fasel, G. Deak, J. Triesch, and J. Movellan, "Combining embodied models and empirical research for understanding the development of shared attention," *Proceedings 2nd International Conference on Development and Learning. ICDL 2002*, pp. 21–27, 2002.

[30] V. V. Hafner and F. Kaplan, "Learning to interpret pointing gestures: Experiments with four-legged autonomous robots," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 3575 LNAI, pp. 225–234, 2005.

[31] M. Doniec, G. Sun, and B. Scassellati, "Active Learning of Joint Attention," *IEEE-RAS International Conference on Humanoid Robots*, pp. 34–39, 2006. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4115577

[32] P. Andry, P. Gaussier, S. Moga, J. Banquet, and J. Nadel, "Learning and Communication in Imitation: An Autnomous Robot Perspective," *IEEE Transaction on Systems, Man and Cybernetics. Part A : Systems and Humans*, vol. 31, no. 5, pp. 431–444, 2001.

[33] S. Boucenna, P. Gaussier, and L. Hafemeister, "Development of joint attention and social referencing," *2011 IEEE International Conference on Development and Learning, ICDL 2011*, pp. 1–6, 2011.

[34] A. D. May, C. Dondrup, and M. Hanheide, "Show me your moves! Conveying navigation intention of a mobile robot to humans," *Mobile Robots (ECMR), 2015 European Conference on*, pp. 1–6, 2015.

[35] H. Ishiguro, T. Ono, M. Imai, and T. Maeda, "Robovie: an Interactive Humanoid Robot," *IEEE International Conference*, vol. 2, no. 1, pp. 1848– 1855, 2002. [Online]. Available: http://onlinelibrary.wiley.com/doi/10.1002/cbdv.200490137/abstract$\delimiter"026E30F$nhttp://www.ingentaconnect.com/content/mcb/049/2001/00000028/00000006/art00006$\delimiter"026E30F$nhttp://ieeexplore.ieee.org/stamp/stamp.jsp?tp={&}arnumber=1014810{&}isnumber=21842

[36] J. F. Ferreira and J. Dias, "Attentional mechanisms for socially interactive robots - A survey," *IEEE Transactions on Autonomous Mental Development*, vol. 6, no. 2, pp. 110–125, 2014.

[37] L. Kelion, "Tesla says autopilot involved in second car crash," jul 2016. [Online]. Available: http://www.bbc.com/news/technology-36783345

[38] S. Thielman, "Fatal crash prompts federal investigation of Tesla self-driving cars," jul 2016. [Online]. Available: https://www.theguardian.com/technology/2016/jul/13/tesla-autopilot-investigation-fatal-crash

[39] J. Vincent, "What counts as artificially intelligent? AI and deep learning, explained," *The Verge*, feb 2016. [Online]. Available: http://www.theverge.com/2016/2/29/11133682/deep-learning-ai-explained-machine-learning

[40] W. Knight, "Can This Man Make AI More Human," *MIT Technology Review*, dec 2015. [Online]. Available: https://www.technologyreview.com/s/544606/can-this-man-make-aimore-human/

[41] B. Goertzel, "Are there Deep Reasons Underlying the Pathologies of Today ' s Deep Learning Algorithms ?" *Artificial General Intelligence Lecture Notes in Computer Science*, vol. 9205, pp. 70–79, 2015.

[42] C. Szegedy, W. Zaremba, and I. Sutskever, "Intriguing properties of neural networks," in *ICLR*, 2014, pp. 1–10. [Online]. Available: http://arxiv.org/abs/1312.6199

[43] A. Nguyen, J. Yosinski, and J. Clune, "Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images," in *CVPR*, 2015. [Online]. Available: http://arxiv.org/abs/1412.1897

[44] (2016, may) The KITTI Vision Benchmark Suite. Online. [Online]. Available: http://www.cvlibs.net/datasets/kitti/

[45] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: An evaluation of the state of the art," *PAMI*, vol. 34, 2012.

[46] K. Fragkiadaki, W. Zhang, G. Zhang, and J. Shi, "Two-Granularity Tracking: Mediating Trajectory and Detection Graphs for Tracking under Occlusions," in *ECCV*, 2012.

[47] S. Y. Bao, M. Bagra, Y.-W. Chao, and S. Savarese, "Semantic Structure from Motion with Points, Regions, and Objects," in *CVPR*, 2012.

[48] (2015) The German Traffic Sign Detection Benchmark. Online. [Online]. Available: http://benchmark.ini.rub.de/?section=gtsdb{&}subsection=dataset

[49] R. Klette. (2014) The .enpeda.. Image Sequence Analysis Test Site (EISATS). Online. [Online]. Available: http://ccv.wordpress.fos.auckland.ac.nz/eisats/{#}1

[50] D. M. Gavrila. (2015) Daimler Pedestrian Benchmark Datasets. [Online]. Available: http://www.gavrila.net/Datasets/Daimler{_}Pedestrian{_}Benchmark{_}D/daimler{_}pedestrian{_}benchmark{_}d.html

[51] UvA Person Tracking from Overlapping Cameras Dataset. [Online]. Available: http://www.gavrila.net/Datasets/Univ{_}{_}of{_}Amsterdam{_}Multi-Cam{_}P/univ{_}{_}of{_}amsterdam{_}multi-cam{_}p.html

[52] 100-Car Naturalistic Driving Study. [Online]. Available: http://www.nhtsa.gov/Research/Human+Factors/Naturalistic+driving+studies:

[53] Transportation Active Safety: 110-Car Naturalistic Driving Study. [Online]. Available: http://www.engr.iupui.edu/{~}yidu/research.html{#}Transportation

[54] SHRP2 Naturalistic Driving Study. [Online]. Available: https://insight.shrp2nds.us/

[55] Virginia Tech Transportation Institute Data Warehouse. [Online]. Available: http://forums.vtti.vt.edu/index.php?/files/category/2-vtti-data-sets/

[56] O. Friard and M. Gamba, "Boris: a free, versatile open-source event-logging software for video/audio coding and live observations," *Methods in Ecology and Evolution*, vol. 7, no. 11, pp. 1325–1330, 2016.